
Subject: Re: SMP experiences with IDL
Posted by [pit](#) on Wed, 25 Feb 1998 08:00:00 GMT
[View Forum Message](#) <> [Reply to Message](#)

In article <34F3190B.41C67EA6@some.place.org>,
Kirt Schaper <xxx@some.place.org> writes:
> Does anyone have first hand experience with IDL (preferably
> on a Linux box) running with multiple processors? Is there
> any speedup? (I'm talking about IDL v5.0 for Unix).

I can only speak for 4.01, but as far as the MP capabilities are
concerned, that won't make a big difference.

1) You won't benefit from having more than 1 processor, as (currently)
there is no threading support. So one IDL job can only make use of
one CPU.

Sounds worse than it is. You can either start 2 different sessions, or
you can do a computation in background and still have the 2nd CPU for
other tasks.

Furtermore, if you have graphic-intensive interactive work, a not too
small part of processing power is needed for the display. This however
is done by the X-Server which is - you guess it - a different process,
and can be run on the other CPU. In Fact I have a lot of
graphics-intensive programs the cause a CPU-usage of 200%: 100 for the
computation and 100 for the display.

2) future versions of IDL *may* support threading, if linked against the
newer c-lib's that are currently under development.

> Our experiences with single processor Pentium/Linux boxes
> suggests that they are at least as fast, if not faster, than
> much more expensive HP, Dec and Sun boxes. Aside from the
> problem of being a little-endian architecture, I haven't been
> able to see the down-side yet.

True. Our Linux PC here (only a P133) is by far the fastest IDL Machine
in the observatory. If you take some care with programming, there's no
problem using the same code and data with different architectures (I
have a Speckle masking program that I can stop in the middle and
continue the work on a different architecture).

> More grist for the Linux performance mill...

0 0
|

_

> Here are some timing results from a simple benchmark program
> (the program simply generated a 100x100x50 random float array
> and convolved it with a 10x10x10 kernel). I know that elapsed
> time is not a very precise benchmark, but the systems were all
> unloaded at the time of the test, and elapsed time is what makes
> a system usable or not.
>
> ; idl version 4.01
> ; SS10/51 (50MHz) ----- elapsed time = 59.1 seconds
> ; Dec 600 5/266 (266MHz) ----- elapsed time = 43.0 seconds
> ; HP 9000 C180 (180MHz) ----- elapsed time = 19.7 seconds
> ; Pentium Pro (200MHz), Linux -- elapsed time = 12.1 seconds
> ; Pentium II (300MHz), Linux --- elapsed time = 9.0 seconds
> ;
> ; idl version 5.0
> ; SS10/51 (50MHz) ----- elapsed time =138.6 seconds
> ; Pentium Pro (200MHz), Linux -- elapsed time = 45.0 seconds
> ; HP 9000 C180 (180MHz) ----- elapsed time = 33.7 seconds
> ; Pentium II (300MHz), Linux --- elapsed time = 31.3 seconds
>
> I find several things interesting about the above experience.
>
> (1) A 200MHz Pentium Pro box is running as fast as a (much more
> expensive, even with 50% academic discount) HP box. This is
> totally contrary to the published SPECfp95_base numbers
> (17.2 for the HP and 5.54 for the Pentium)

I also find the comparison PPro/PII very interesting. It also shows that
Linux makes superior use of the PPro's true 32-bit branching prediction
etc. which are severely reduced for the PII (to be better
Win-compatible): Modulo the MHz, the "next generation Chip" is no bit
faster than the old one...

> (2) RSI did something quite bad to the convolution function
> from v4 to v5.

One reason more to stay with 4.01

Peter

--

~~~~~  
~~~~~

Peter "Pit" Suetterlin <http://www.uni-sw.gwdg.de/~pit>
Universitaets-Sternwarte Goettingen

Tel.: +49 551 39-5048

pit@uni-sw.gwdg.de

-- * -- * ...-- * -- * ...-- * -- * ...-- * -- * ...-- * -- * ...-- * --

Come and see the stars!

<http://www.kis.uni-freiburg.de/~ps/SFB>

Sternfreunde Breisgau e.V.

Tel.: +49 7641 3492